



Cold
Spring
Harbor
Laboratory

Introduction to cloud computing

Kelsy Cotto, Felicia Gomez, Obi Griffith, Malachi Griffith,
My Hoang, Chris Miller, Huiming Xia


Advanced Sequencing Technologies & Bioinformatics Analysis November 6-20, 2022



Learning Objectives

- Introduction to cloud computing concepts
- Introduction to cloud computing providers
- Use the Amazon EC2 console to create an instance for each student
 - Will be used for many hands-on tutorials throughout the course
- How to log into your cloud instance

Disk Capacity vs Sequencing Capacity, 1990-2012



About DNA and computers

- We hit the \$1000 genome* in ~2016
 - Need to think about the \$100 genome
- The doubling time of sequencing has been ~5-6 months.
- The doubling time of storage and network bandwidth is ~12 months.
- The doubling time of CPU speed is ~18 months.
- The cost of sequencing a base pair will eventually equal the cost of storing a base pair

What is the general biomedical scientist to do?

- Lots of data
- Poor IT infrastructure in many labs
- Where do they go?
- Get bigger hardware?
- Write more grants?

Cloud computing providers

- Amazon AWS
 - <https://aws.amazon.com/>
- Google cloud
 - <https://cloud.google.com/>
- Microsoft Azure
 - <https://azure.microsoft.com/en-us/>
- More...

Amazon Web Services (AWS)

- Infinite storage (scalable): S3 (simple storage service)
- Compute per hour: EC2 (elastic cloud computing)
- Ready when you are High Performance Computing
- Multiple football fields of HPC throughout the world



Some of the challenges of cloud computing:

- Not cheap
- Getting files to and from there
- Standardization can be a challenge if you don't control hardware
- PHI: personal health information & security concerns
 - In the USA: HIPAA act, PSQIA act, HITECH act, Patriot act, CLIA and CAP programs, etc.
 - <http://www.biostars.org/p/70204/>

Some of the advantages of cloud computing:

- There are better ways of transferring large files, and now AWS makes it free to upload files.
- A number of datasets exist on AWS (e.g. 1000 genome data, TCGA).
- Many useful bioinformatics AMI's (Amazon Machine Images) exist on AWS: e.g. cloudbiolinux & CloudMan (Galaxy) – now one for this course!
- Many flavors of cloud available, not just AWS

Key AWS concepts and terminology

- **AWS** - Amazon Web Services. A collection of cloud computing services provided by Amazon.
- **EC2** - Elastic Compute. An AWS service that allows you to configure and rent computers to meet your compute needs on an as needed basis.
- **EBS** - Elastic Block Storage. A data storage solution that allows you to rent disk storage and associate that storage with your compute resources. EBS volumes are generally backed by SSD devices.

Key AWS concepts and terminology

- **S3** - Simple storage service. Cheaper than EBS and allows for storage of larger amounts of data with some drawbacks compared to EBS. S3 volumes store data as objects that are accessed by an API or command line interface or other application designed to work with S3. EBS volumes on the other hand can be mounted as if they were a local disk drive associated with the Instance.
- **SSD** - Solid state drive. A particular type of storage hardware that is generally faster and more expensive than traditional hard drives.

What is difference between the 'Start', 'Stop', 'Reboot', and 'Terminate' (Instance States)?

- Start – turn on an EC2 instance that you have previously created
- Stop – turn off an EC2 instance that you have previously created
- Reboot – restart an EC2 instance
- Terminate – permanently stop and destroy an EC2 instance. Any associated EBS volumes may also be destroyed at this time depending on configuration

What is an AMI/snapshot?

- AMI (Amazon Machine Image) – a template that specifies how to launch EC2 instances
 - Root volume with operating system (OS), pre-installed applications, etc
 - Launch permissions that determine who can use the AMI
 - Specification of (data) volumes to attach when launched
- You can create an AMI for any instance you have created/configured
- AMI can be made public for sharing (region-specific)
- Creating an AMI involves creating a snapshot of the root and any attached volumes. You will be charged to store this snapshot.

I can not log into my EC2 instance, what might have gone wrong?

- Is your instance running?
- Are you providing the correct path to your key file?
- Is it the correct key file?
- Have you set the permissions for your key file correctly?
- Did you specify a valid user for your AMI (e.g., ubuntu)?
- Did you specify the correct IP address?
- Does the Security Group for the instance allow access for your connection protocol (e.g., SSH) and location?

How much does it cost to use AWS EC2 resources?

Viewing 489 of 489 available instances

< 1 2 3 4 5 6 7 ... 25 >

Instance name ▲	On-Demand hourly rate ▼	vCPU ▼	Memory ▼	Storage ▼	Network performance ▼
a1.medium	\$0.0255	1	2 GiB	EBS Only	Up to 10 Gigabit
a1.large	\$0.051	2	4 GiB	EBS Only	Up to 10 Gigabit
a1.xlarge	\$0.102	4	8 GiB	EBS Only	Up to 10 Gigabit
a1.2xlarge	\$0.204	8	16 GiB	EBS Only	Up to 10 Gigabit
a1.4xlarge	\$0.408	16	32 GiB	EBS Only	Up to 10 Gigabit
a1.metal	\$0.408	16	32 GiB	EBS Only	Up to 10 Gigabit
t4g.nano	\$0.0042	2	0.5 GiB	EBS Only	Up to 5 Gigabit
t4g.micro	\$0.0084	2	1 GiB	EBS Only	Up to 5 Gigabit
t4g.small	\$0.0168	2	2 GiB	EBS Only	Up to 5 Gigabit
t4g.medium	\$0.0336	2	4 GiB	EBS Only	Up to 5 Gigabit
t4g.large	\$0.0672	2	8 GiB	EBS Only	Up to 5 Gigabit
t4g.xlarge	\$0.1344	4	16 GiB	EBS Only	Up to 5 Gigabit

Data transfer (GB): In: free; Out: first 100 GB free, \$0.05-0.09 per GB depending on amount

EBS storage (GB/Month): \$0.08

S3 storage (GB/Month): \$0.023 standard, \$0.0125 infrequent access, or \$0.004 glacier

Why am I still getting a monthly bill?

- Generally you get an accounting of usage and cost on a 30 day cycle
 - Pricing is per instance-second consumed for each instance type.
 - Also charges for storage, transfers, etc
- Be aware of regions!
- Even when an instance is stopped, storage for root or other EBS volumes persist
- Creating AMIs/snapshots requires storage
- Explore the billing and cost management tools of AWS to track your spending, set warnings, etc

Amazon AWS documentation

https://rnabio.org/module-00-setup/0000/06/01/Intro_to_AWS/

<http://aws.amazon.com/console/>

In this workshop:

- Some tools (data) are
 - on your computer
 - on the web
 - on the cloud.
- You will become efficient at traversing these various spaces, and finding resources you need, and using what is best for you.
- There are different ways of using the cloud:
 1. Command line (like your own very powerful Unix box)
 2. With a web-browser (e.g. Galaxy): not in this workshop

Things we have set up:

- Loaded data files to a web server (genomedata.org)
- We started with a base Ubuntu 20.04 (Linux) instance and loaded a whole bunch of software for NGS analysis.
- We will copy this and create separate instances for everybody in the class.
- We've simplified the security: you basically all have the same login and file access, and opened ports. In your own world, you would be more secure to protect your data.

Logging into Amazon AWS

Go to course wiki, “Log into AWS” page



Log into AWS

[« Introduction to AWS](#)

[Course](#)

[Unix »](#)

Using cloud computing to complete this course involves two major components: (1) Launching an instance on the cloud (essentially renting a virtual computer by the hour and turning it on) and (2) logging into that instance).

Covered in this section: logging into AWS EC2 console, starting an instance from the course AMI, configuring it in the console (select instance AMI, instance type, instance details, storage volumes, tags, security group, and key pairs).



https://rnabio.org/module-00-setup/0000/07/01/Log_into_AWS/

Login to AWS console



Account ID or alias

IAM user name

Password

Sign In



[Sign-in using root account credentials](#)

[Forgot password?](#)

<https://cshlworkshops.signin.aws.amazon.com/console>

Select "EC2" service

The screenshot displays the AWS Management Console interface. At the top, the search bar contains the text "EC2". A red arrow points to this search bar. Below the search bar, a white box with the text "Search for EC2" is overlaid. The search results are categorized into "Services" and "Features". Under "Services", the first result is "EC2" (Virtual Servers in the Cloud), which is highlighted with a blue border. Other services listed include "EC2 Image Builder", "AWS Firewall Manager", and "GuardDuty". Under "Features", there are two results for "Dashboard" and "Limits", both labeled as "EC2 feature". On the right side of the console, the region is set to "N. Virginia", and the user is identified as "kelsycotto @ 5772-5572-5291". A red arrow points to the region dropdown menu. A large text overlay on the right side of the console reads "Make sure you are in N. Virginia region".

From EC2 Dashboard, launch a new Instance

The screenshot shows the AWS Management Console interface for the EC2 Dashboard. The top navigation bar includes the AWS logo, 'Services', a search bar, and the account name 'kelsycotto @ 5772-5572-5291' in the 'N. Virginia' region. The left-hand navigation pane lists various EC2-related options, with 'EC2 Dashboard' highlighted and a red arrow pointing to it. Below this, categories like 'Instances', 'Images', and 'Elastic Block Store' are visible. The main content area is divided into several sections: 'Resources' showing a summary of EC2 resources in the US East (N. Virginia) Region (e.g., 2 running instances, 4 volumes); 'Launch instance' with a prominent orange 'Launch instance' button and a red arrow pointing to it; 'Service health' indicating that the service is operating normally; and 'Scheduled events' for the region. On the right side, there are panels for 'Account attributes' (including supported platforms and VPC settings) and 'Explore AWS' (with links to GuardDuty and cost reduction tips).

Name your instance “FirstnameLastname” (e.g. KelsyCotto)

EC2 > Instances > Launch an instance

Launch an instance [Info](#)

Amazon EC2 allows you to create virtual machines, or instances, that run on the AWS Cloud. Quickly get started by following the simple steps below.

Name and tags [Info](#)

Name

 [Add additional tags](#)

Application and OS Images (Amazon Machine Image) [Info](#)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. Search or Browse for AMIs if you don't see what you are looking for below

- Amazon Linux
- macOS
- Ubuntu
- Windows
- Red Hat

[Browse more AMIs](#)

Including AMIs from AWS, Marketplace and the Community

Summary

Number of instances [Info](#)

Software Image (AMI)

Amazon Linux 2 Kernel 5.10 AMI...[read more](#)
ami-09d3b3274b6c5d4aa

Virtual server type (instance type)

t1.micro

Firewall (security group)

New security group

Storage (volumes)

1 volume(s) - 8 GiB

Cancel [Launch instance](#)

Important: Don't forget to name your instance!
(FirstnameLastname)

Choose an AMI – Find the CSHL SEQTEC 2022 AMI in the My AMIs

EC2 > Instances > Launch an instance

Launch an instance [Info](#)

Amazon EC2 allows you to create virtual machines, or instances, that run on the AWS Cloud. Quickly get started by following the simple steps below.

Name and tags [Info](#)

Name: [Add additional tags](#)

Application and OS Images (Amazon Machine Image) [Info](#)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. Search or Browse for AMIs if you don't see what you are looking for below.

Search our full catalog including 1000s of application and OS images

Recents | **My AMIs** | [All AMIs](#)

Owned by me Shared with me

[Browse more AMIs](#)
Including AMIs from AWS, Marketplace and the Community

Amazon Machine Image (AMI)

cshl-seqtec-2022 ami-03fd53fd1063112d2 2022-11-08T16:17:23.000Z Virtualization: hvm ENA enabled: true Root device type: ebs

Summary

Number of instances [Info](#):

Software Image (AMI) (highlighted in red box):
cshl-seqtec-2022
ami-03fd53fd1063112d2

Virtual server type (instance type): t2.micro

Firewall (security group): New security group

Storage (volumes): 2 volume(s) - 282 GiB

Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GiB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet.

[Cancel](#) [Launch instance](#)

Choose "m5.2xlarge" instance type, then "Next: Configure Instance Details".

Amazon Machine Image (AMI)

csbl-seqtec-2022
ami-03fd53fd1063112d2
2022-11-08T16:17:23.000Z Virtualization: hvm ENA enabled: true Root device type: ebs

Description

-

Architecture	AMI ID
x86_64	ami-03fd53fd1063112d2

Software Image (AMI)

csbl-seqtec-2022
ami-03fd53fd1063112d2

Virtual server type (instance type)
m5.2xlarge

Firewall (security group)

New security group

Storage (volumes)

2 volume(s) - 282 GiB

Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GiB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet.

Instance type [Info](#)

Instance type

m5.2xlarge
Family: m5 8 vCPU 32 GiB Memory
On-Demand Linux pricing: 0.384 USD per Hour
On-Demand Windows pricing: 0.752 USD per Hour

[Compare instance types](#)

Choose an existing key pair: "cshl_2022_student"

▼ **Instance type** [Info](#)

Instance type

m5.2xlarge
Family: m5 8 vCPU 32 GiB Memory
On-Demand Linux pricing: 0.384 USD per Hour
On-Demand Windows pricing: 0.752 USD per Hour

[Compare instance types](#)


▼ **Key pair (login)** [Info](#)

You can use a key pair to securely connect to your instance. Ensure that you have access to the selected key pair before you launch the instance.

Key pair name - *required*

cshl_2022_student

[Create new key pair](#)



Select an Existing Security Group, choose "SSH_HTTP". Then hit "Review and Launch".

The screenshot displays the AWS console interface for configuring an instance. On the left, the 'Network settings' section is expanded, showing options for Network (vpc-ad2c8fd7), Subnet (No preference), and Auto-assign public IP (Enable). The 'Firewall (security groups)' section is active, with two radio buttons: 'Create security group' (unselected) and 'Select existing security group' (selected, indicated by a red arrow). Below these is a dropdown menu labeled 'Select security groups' with a red arrow pointing to the selected item 'SSH/HTTP/Jupyter sg-0fd7356e013ca81a3'. A 'Compare security group rules' button is also visible. On the right, the 'Software Image (AMI)' is 'cshl-seqtec-2022', the 'Virtual server type (instance type)' is 'm5.2xlarge', and the 'Firewall (security group)' is 'SSH/HTTP/Jupyter' (highlighted with a red box). A 'Storage (volumes)' section shows '2 volume(s) - 282 GiB'. A 'Free tier' information box is also present.

You should see 1 32 GiB root volume and 1 250 GiB EBS volume as the two storage volumes.

▼ **Configure storage** [Info](#) Advanced

1x GiB ▼ Root volume (Not encrypted)

1x GiB ▼ EBS volume (Not encrypted) Remove

i Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage ✕

Add new volume

The selected AMI contains more instance store volumes than the instance allows. Only the first 0 instance store volumes from the AMI will be accessible from the instance

0 x File systems Edit

Select “Enable” for termination protection.

The screenshot displays the AWS console configuration page for a new EC2 instance. The left pane is titled "Advanced details" and contains several sections:

- Purchasing option:** "Request Spot Instances" is unchecked.
- Domain join directory:** A dropdown menu is set to "Select".
- IAM instance profile:** A dropdown menu is set to "Select".
- Hostname type:** A dropdown menu is set to "IP name".
- DNS Hostname:** Three checkboxes are checked: "Enable IP name IPv4 (A record) DNS requests", "Enable resource-based IPv4 (A record) DNS requests", and "Enable resource-based IPv6 (AAAA record) DNS requests".
- Instance auto-recovery:** A dropdown menu is set to "Select".
- Shutdown behavior:** A dropdown menu is set to "Stop".
- Stop - Hibernate behavior:** A dropdown menu is set to "Select".
- Termination protection:** A dropdown menu is set to "Enable". A large red arrow points to this dropdown.

The right pane is titled "Summary" and shows the following configuration:

- Number of instances:** 1
- Software Image (AMI):** cshl-seqtec-2022 (ami-03fd53fd1063112d2)
- Virtual server type (instance type):** m5.2xlarge
- Firewall (security group):** SSH/HTTP/Jupyter
- Storage (volumes):** 2 volume(s) - 282 GiB

A "Free tier" notification box is visible, stating: "Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GiB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet." At the bottom of the console, there are "Cancel" and "Launch instance" buttons.

Review the details of your instance and hit Launch

▼ Summary

Number of instances [Info](#)

Software Image (AMI)
cshl-seqtec-2022
ami-03fd53fd1063112d2

Virtual server type (instance type)
m5.2xlarge

Firewall (security group)
SSH/HTTP/Jupyter

Storage (volumes)
2 volume(s) - 282 GiB

Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GiB of EBS storage, 2 million IOs, 1 GB of snapshots, and 100 GB of bandwidth to the internet. ×

Cancel Launch instance

View Instances to see your new instance spinning up!

EC2 > Instances > Launch an instance



Success

Successfully initiated launch of instance (i-00e1dc9cd6843478d)

▶ [Launch log](#)

Next Steps

Create billing and free tier usage alerts

To manage costs and avoid surprise bills, set up email notifications for billing and free tier usage thresholds.

[Create billing alerts](#) 

Connect to your instance

Once your instance is running, log into it from your local computer.

[Connect to instance](#) 

[Learn more](#) 

Connect an RDS database New

Configure the connection between an EC2 instance and a database to allow traffic flow between them.

[Connect an RDS database](#) 

[Create a new RDS database](#)  [Learn more](#) 

[View all instances](#)

Find YOUR instance, select it, and then hit connect for instructions on how to connect (It may take some time for your instance to be ready)

The screenshot displays the AWS Management Console interface for EC2 instances. On the left, a navigation sidebar lists various services like EC2 Dashboard, Events, Tags, Limits, and Images. The main area shows a table of instances with columns for Name, Instance ID, Instance state, Instance type, Status check, Alarm status, Availability Zone, Public IPv4 DNS, Public IPv4 address, and Elastic IP. The instance 'KelsyCotto' is selected, and its details are shown below. A red box highlights the 'Connect' button in the top right, and a red arrow points to the 'Connect' button. Another red arrow points to the 'Public IPv4 address' field in the instance details, which is '100.26.22.19'.

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv4 ...	Elast
Kcotto	i-0302a728fbd4095f	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-234-50-3.comp...	18.234.50.3	-
Instructor_test	i-088f1b7da8b8d0656	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-212-77-248.co...	18.212.77.248	-
KelsyCotto	i-00e1dc9cd6843478d	Running	m5.2xlarge	Initializing	No alarms	us-east-1a	ec2-100-26-22-19.com...	100.26.22.19	-

Instance: i-00e1dc9cd6843478d (KelsyCotto)

Details | Security | Networking | Storage | Status checks | Monitoring | Tags

Instance summary Info

Instance ID i-00e1dc9cd6843478d (KelsyCotto)	Public IPv4 address 100.26.22.19 open address	Private IPv4 addresses 172.31.28.82
IPv6 address -	Instance state Running	Public IPv4 DNS ec2-100-26-22-19.compute-1.amazonaws.com open address
Hostname type IP name: ip-172-31-28-82.ec2.internal	Private IP DNS name (IPv4 only) ip-172-31-28-82.ec2.internal	Elastic IP addresses -
Answer private resource DNS name IPv4 (A)	Instance type m5.2xlarge	

Take note of your Public DNS/IP and the instructions on changing permissions for the key file (Note, we will login as ubuntu NOT root)

EC2 > Instances > i-00e1dc9cd6843478d > Connect to instance

Connect to instance [Info](#)

Connect to your instance i-00e1dc9cd6843478d (KelsyCotto) using any of these options

[EC2 Instance Connect](#) | [Session Manager](#) | **[SSH client](#)** | [Elastic console](#)

Instance ID
i-00e1dc9cd6843478d (KelsyCotto)

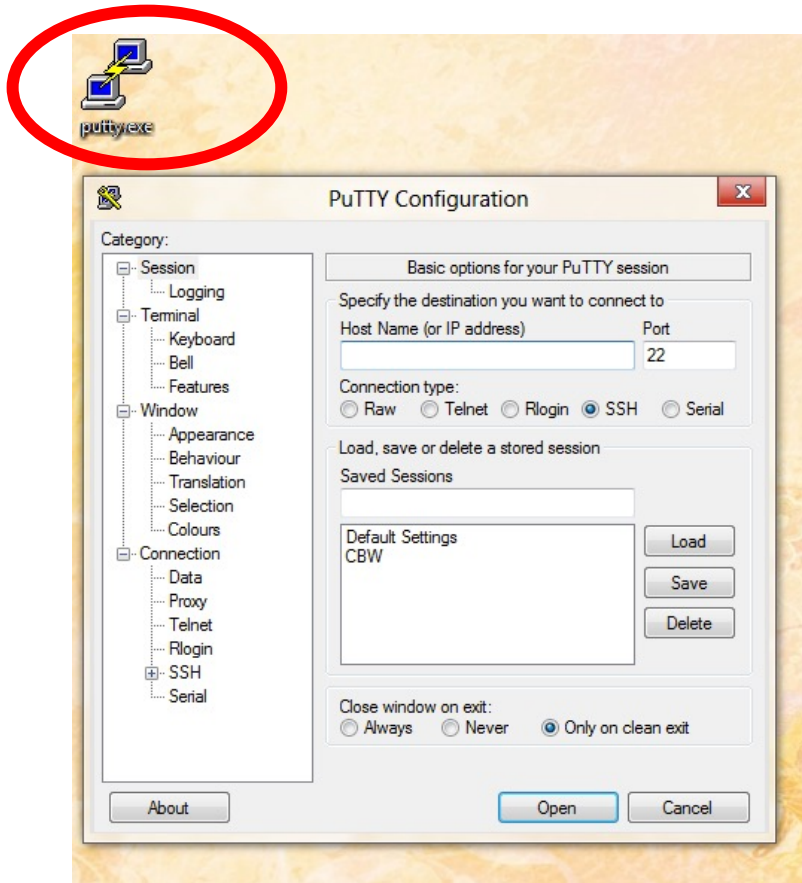
1. Open an SSH client.
2. Locate your private key file. The key used to launch this instance is cshl_2022_student.pem
3. Run this command, if necessary, to ensure your key is not publicly viewable.
❏ `chmod 400 cshl_2022_student.pem`
4. Connect to your instance using its Public DNS:
❏ `ec2-100-26-22-19.compute-1.amazonaws.com`

Example:
❏ `ssh -i "cshl_2022_student.pem" root@ec2-100-26-22-19.compute-1.amazonaws.com`

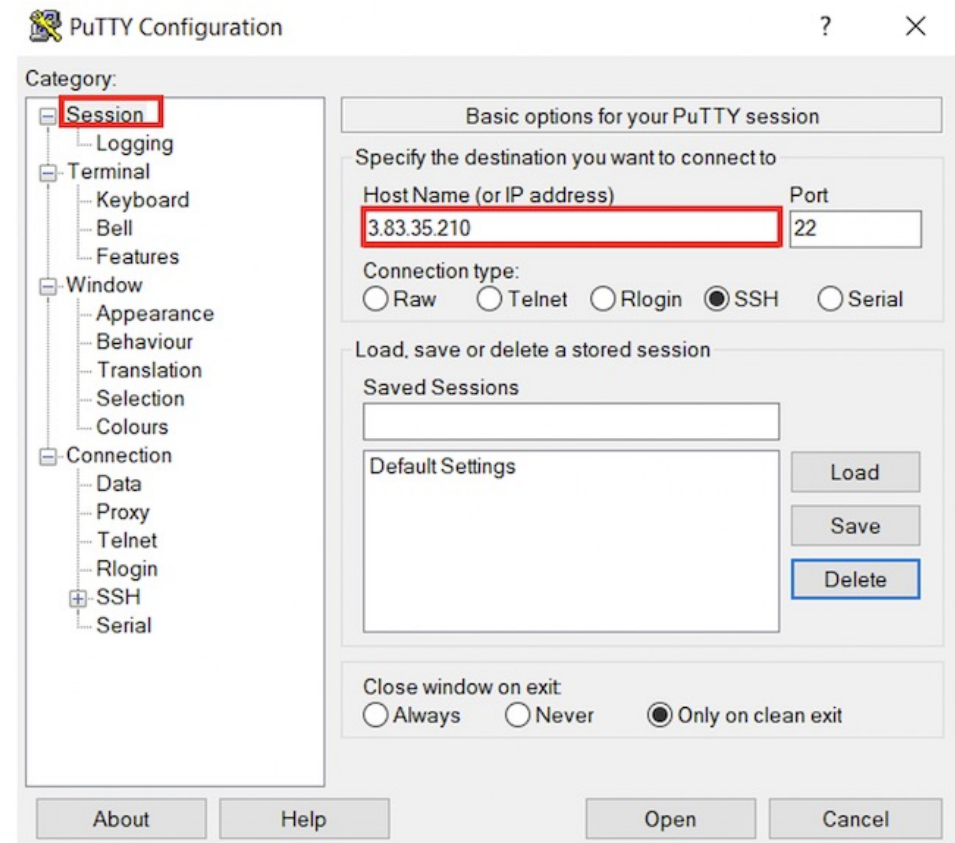
Note: In most cases, the guessed user name is correct. However, read your AMI usage instructions to check if the AMI owner has changed the default AMI user name.

Logging into your instance (Windows)

Open PuTTY

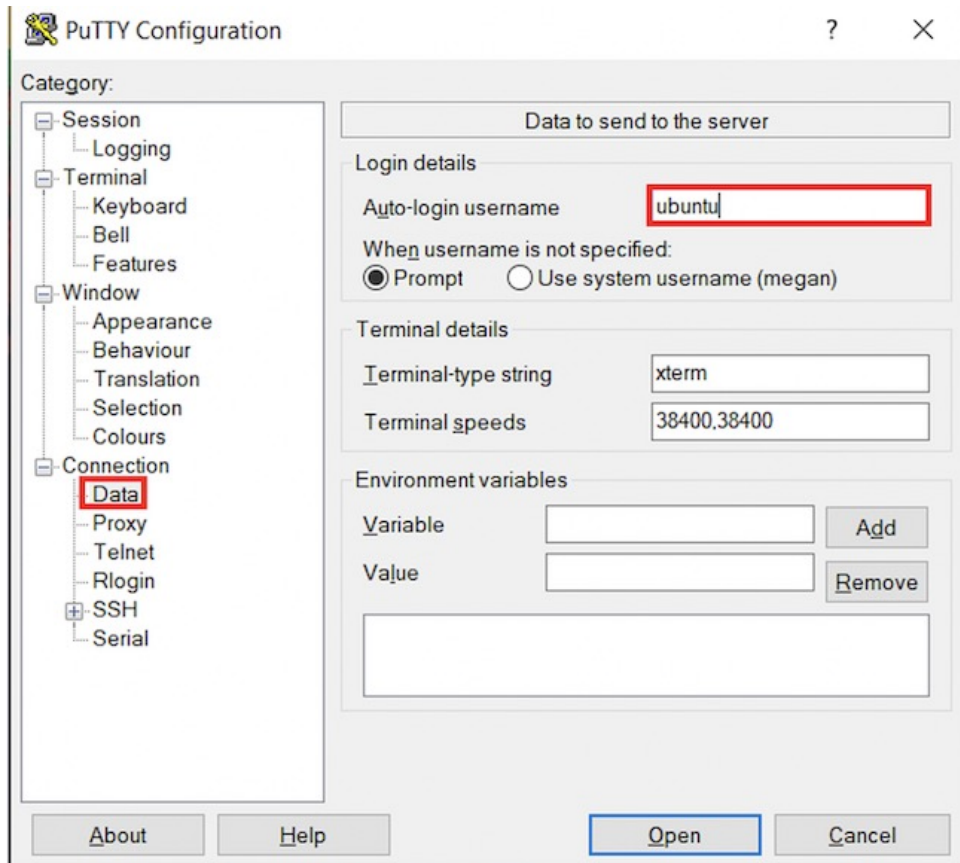


Enter the Host Name (IP address)

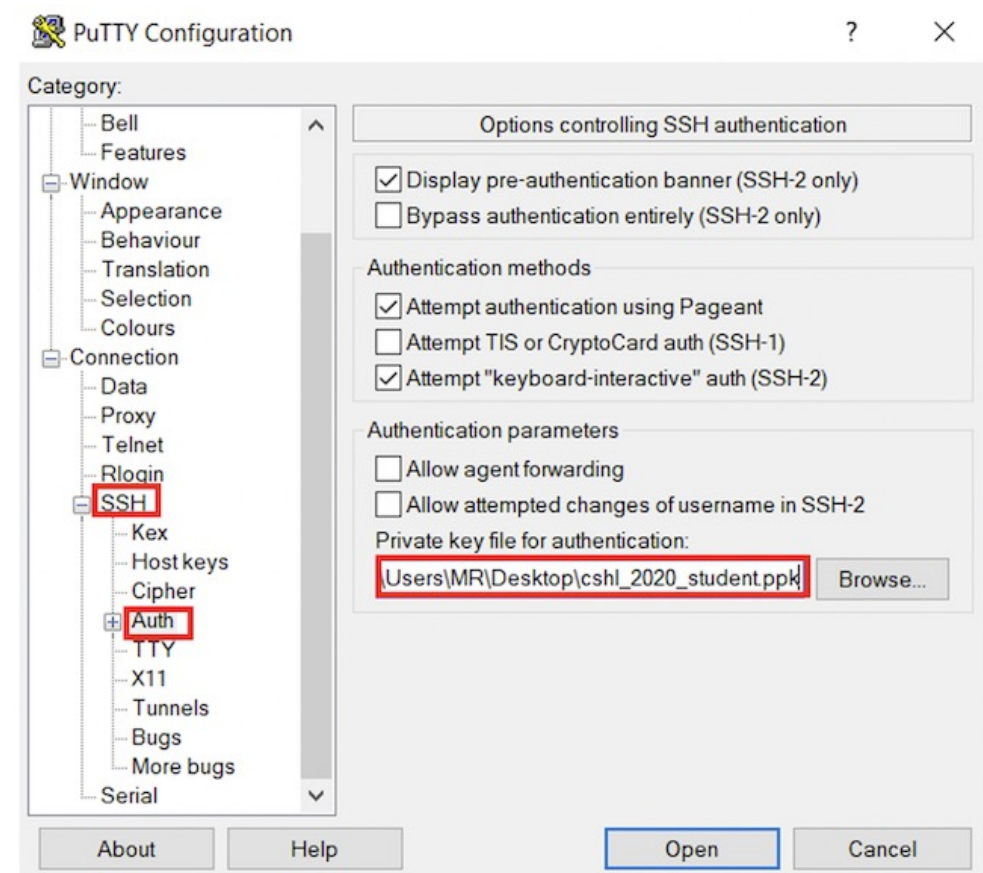


Logging into your instance (Windows)

Choose Connection -> Data
Enter the username 'ubuntu'

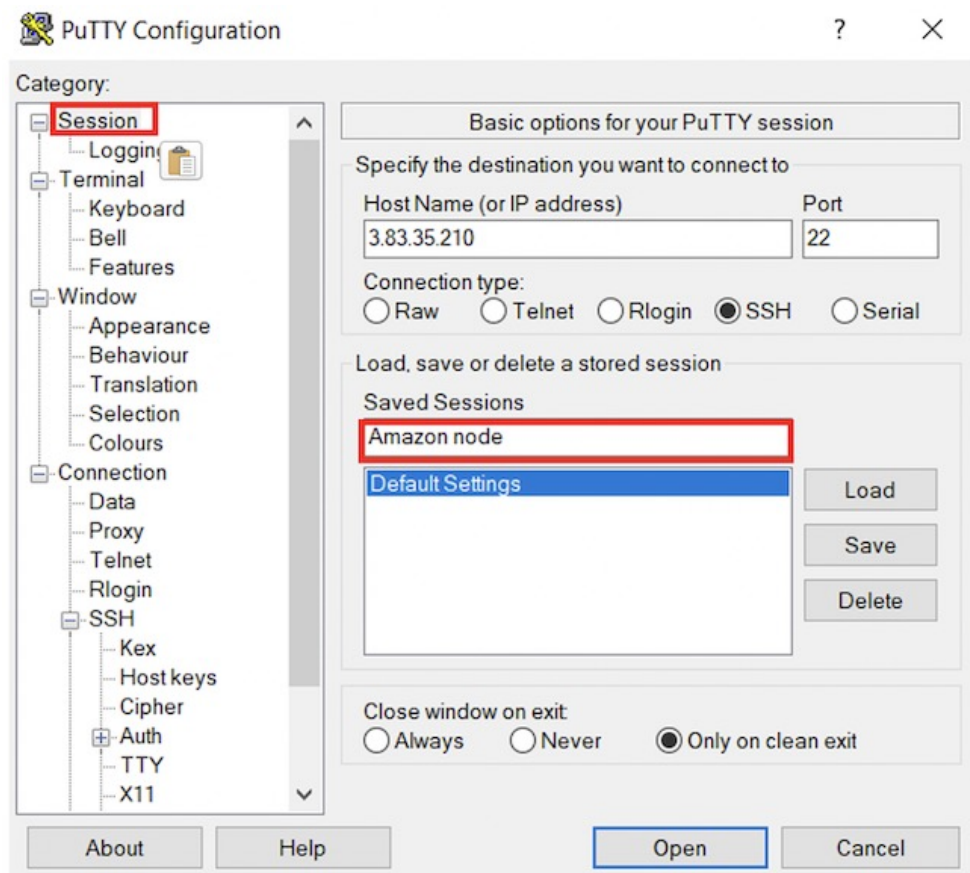


Choose SSH -> Auth
Browse to Private key (ppk) file

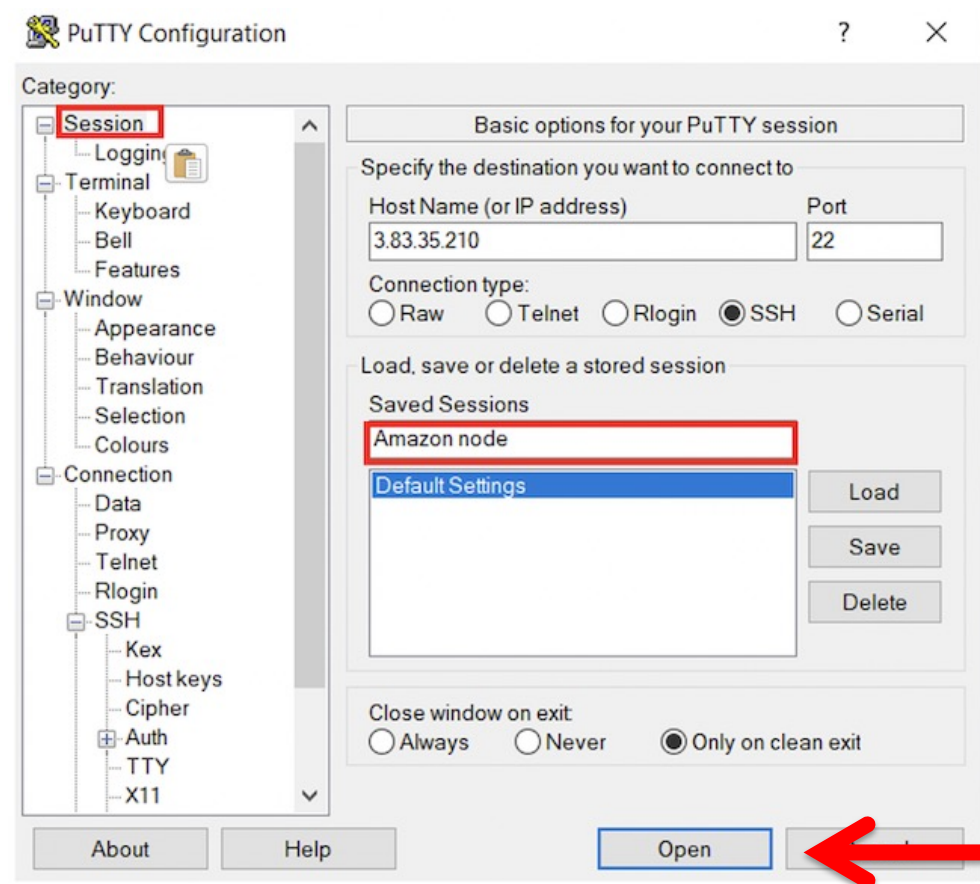


Logging into your instance (Windows)

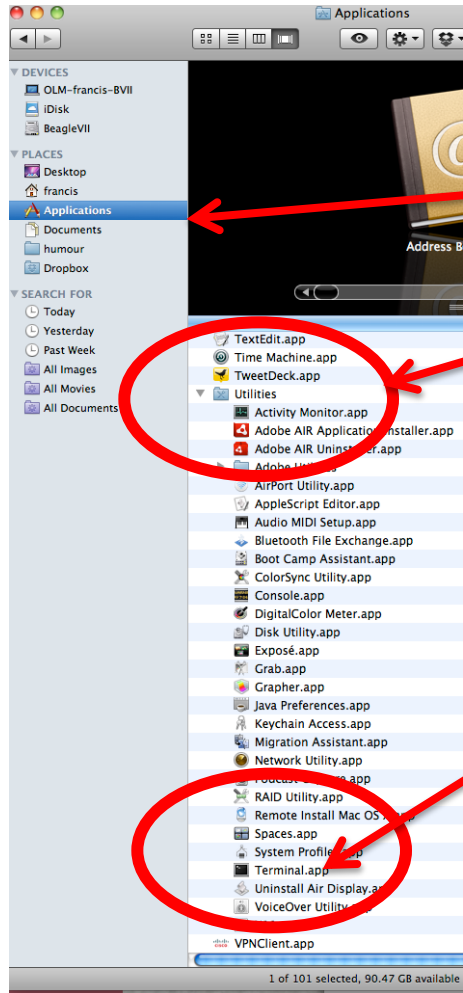
Choose Session
Save your session as "Amazon Node"



Double-click saved "Amazon Node" session OR
Select "Amazon Node" session and click Open



Logging into your instance (Mac)

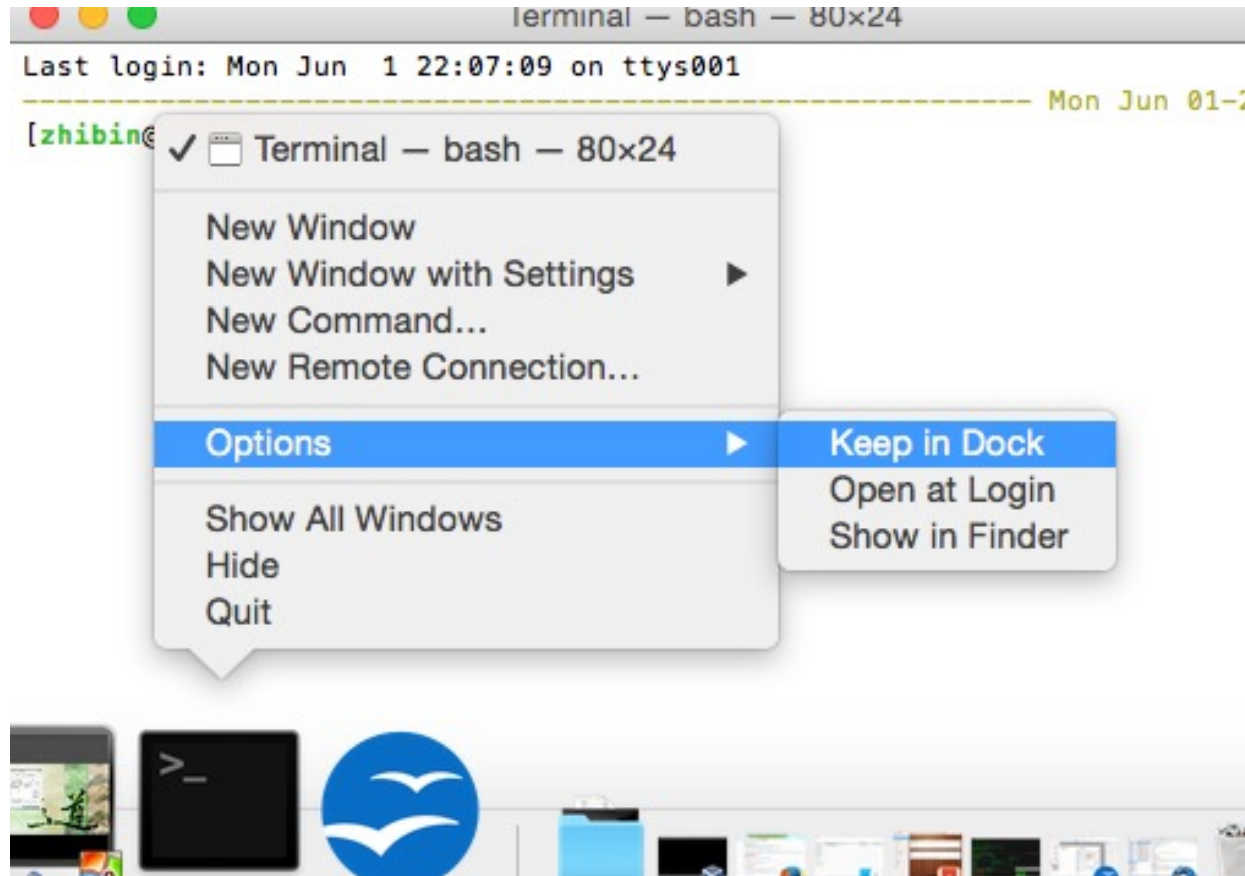


In a Finder window
'Applications' -> 'Utilities' -> 'Terminal'

Or on your dock



Add the terminal App to your dock



Creating a working directory on your Mac called 'cshl'

```
obis-air:~ ogriffit$ pwd
/Users/ogriffit
obis-air:~ ogriffit$ ls
Applications      Desktop           Dropbox           Movies            Public            gittemp          temp
Attachments       Documents        Google Drive     Music             bin               igv
Box Sync          Downloads        Library           Pictures          git               ncbi
obis-air:~ ogriffit$ mkdir cshl
obis-air:~ ogriffit$ cd cshl
obis-air:cshl ogriffit$ ls -la
total 0
drwxr-xr-x  2 ogriffit  staff   68 Nov 13 22:18 .
drwxr-xr-x+ 58 ogriffit  staff  1972 Nov 13 22:18 ..
obis-air:cshl ogriffit$
```

```
mkdir cshl
cd cshl
```

Obtain the course SSH key file

- NOTE for Mac users. You will need to use a “.pem” file
- **NOTE for Windows Users.** You will need to use a “.ppk” file instead.
 - This is created from the “.pem” file.
 - <https://aws.amazon.com/premiumsupport/knowledge-center/convert-pem-file-into-ppk/>
- The SSH key file will be used to securely login to your student instance on the cloud

Save the pem/ppk file you received via email/slack to your new cshl folder

Viewing the 'key' file once downloaded

```
cat cshl_2021_student.pem
```

```
Kelsys-MacBook-Pro:keys kcotto$ cat cshl_2019_student.pem
-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAACAQEAgtGtpnqERtEu/SCmeF2r1HMESMao0fEJiAwQwk2/SNXK8izr0IH0zTVvmE1D
VUwWq7pkvhhJh05pDb+2U9HiZe3sXlv3S1NrkATYF/NsrpwB+q1vwqzGW9sQ6uj45RWrpkiZlsaj
TQZmyFRu+tLJTRU3hQDqA0MRWTx1Wxv0GfZuZy/qb+DALuFQsInrEKnijrwdLmd6usaBTvhc0gFS
B9oEelH0bZHJTzFw/wP+Z0uZq0Ujir7Qw0LTM45QH/L0dBdU13k/mBeez00yvnKMwj8E4Xi0rQ0t
hHtQ7F9iSILK80W1rRH0qwxwt9ycEH1JtNMQmUtif0vE2XJ6L06chQIDAQABAoIBABG7P/FHu/Qp
WFgg+89myuqR6GvA2X55CFSzZfYg0aQyrj5jDleFtdu2uXiISG8gUBZYvlzxx82a0C0P5j04SBq0
xD/qRlukY/jyXyPn77w/ExmaNoLJj1W9RUSH0JYLIZvFPGes3u5zGSGDTSNDhw3sSdWhq1FX3L
7vY5b6UAQgahXf0dpGFxt6P6qb/BKFQfSthXk7GXMyS/kr4w7Z1hPWHrMSu2UBdS/a/beAVN76qU
E/10HR30oAuqghusZabpbX7MJl00VcviAQgeF8Z+xf2uugDEAKut1PxW0+yvGM4SpZZ0skFZz6YI
pBnX4ELWPhCeMzq4Ml8QY6ptR2UCgYEAuc6Yfnlymg24NJU8PANx8navTi50WYuXWxrj0mrLr3SR
/XY+w26cgipM+K5eQfXsr4Yb8BQKjRktMzBzf5nKdaX4pzYIquQH02B0HDhTooHAhkbTMTmKukv4
oEW06wcEE6RQi fw4xbMEnfQfHJBI21am+jwI8Xb7idwMG4pU/nsCgYEAtD0x4bNC1X3A5by50uY0
WXrtgQszCXYcbkrpbjRET12f9hgZ9MRMHY/xH/XGvMutZSFV2rCZRwd7lm+QWGadk/MQS0kouzW6
gSasyjFq+MKCKqYnS3/JTbx8yrLZmz10LtX6pwwmg0Zy8aJjYHo9a2/EI8Tjh2d0SxeadIRVYP8C
gYAToiXww1Vdu+dj/7TDLqYctdH0VaxJX/utI9QQ3yoIryuh+bWmFvEIvAmIGXyyQZRyoZwgIS4A
PNH03+bEa+69wbzlhksiK5g8GKgISVdlC4rZZXB5ehgTmWV7Igj89y/SF4G/Ityo30K0ohALh597
NcvNEzzqruTja1IIMvTKMwKBgQCkM+QP1Tqc0TbVlFvClviXuJBLsiJLClmYeZL0nZVmIMusbhxX
b8ZQYGSyUz09nuLXaulGLQDvXvf089CzWLLSomxBoHlFJQvGwa9FfYQRIVPHuqut8rs4oPGn0QzC
h7M7QCJcr00oAcrSLLkQmgz+phIw7BzFr039J4HFIRInjQKBgQCtdEvcbytk8Jh4WH3z0wpkc43f
U8DZhZwjRQpGwLD8CPj9RgRnE4+1PCH6s/RLQf7SiE1ZjX/0Ud0WPEvr0j5sVjy0IujohRbty0CM
oqWeSeUb1sLogRvMrTfCEpl/rz3GpoQLSC/5s6XvjnnKK8RN8s7MseLuuJ63T/CRBpIs8A==
-----END RSA PRIVATE KEY-----Kelsys-MacBook-Pro:keys kcotto$
```

Changing file permissions of your 'key' file (Mac/Linux)

ls -l (long listing)

```
-rw-r--r--@ 1 kcotto staff 1696 Nov 9 09:19 cshl_2022_student.pem
```

```
rwX : owner
```

```
  rwX : group
```

```
    rwX: world
```

```
r read  (4)
```

```
w write (2)
```

```
x execute (1)
```

Which ever way you add these 3 numbers, you know which integers were used (6 is always 4+2, 5 is 4+1, 4 is by itself, 0 is none of them etc ...)

So, when you have:

chmod 400 <file name>

It is "r" for the the file owner **only**






Logging into your instance

Mac/Linux

```
cd ~/cshl  
chmod 400 cshl_2022_student.pem  
ssh -i cshl_2022_student.pem ubuntu@[YOUR PUBLIC IP]
```

Copying files from AWS to your computer (using a web browser)

Index of /

<u>Name</u>	<u>Last modified</u>	<u>Size</u>	<u>Description</u>
 anaconda3/	2022-11-07 16:56	-	
 bin/	2022-11-08 10:39	-	
 src/	2022-11-08 10:25	-	
 tmp	2022-07-04 07:24	736K	
 workspace/	2022-11-08 10:58	-	

Apache/2.4.41 (Ubuntu) Server at 100.26.22.19 Port 80

http://[YOUR PUBLIC DNS OR IP]/

Logging out of your instance

Mac/Linux – simply type exit

exit

Note, this disconnects the terminal session (ssh connection) to your cloud instance. But, your cloud instance is still running! See next slide for how to stop your instance.

When you are done for the day you can “Stop” your instance – Don’t Terminate!

The screenshot shows the AWS Management Console's EC2 Instances page. The left sidebar contains navigation options like 'EC2 Dashboard', 'Instances', and 'Images'. The main area displays a table of instances. The 'KelsyCotto' instance is selected, and a context menu is open over it, with the 'Stop instance' option highlighted. Red arrows point to the 'Instances' tab in the left sidebar, the 'KelsyCotto' instance in the table, and the 'Stop instance' option in the menu.

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv4 ...	Elast
Kcotto	i-0302a728fbd4095f	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-234-50-3.comp...	18.234.50.3	-
Instructor_test	i-088f1b7da8b8d0656	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-212-77-248.co...	18.212.77.248	-
KelsyCotto	i-00e1dc9cd6843478d	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-54-221-110-174.co...	54.221.110.174	-

Go to AWS EC2 Dashboard, select “Instances” tab, then find your instance. Right-click and chose ‘Stop instance’

Next morning, you can “Start” your instance again

The screenshot displays the AWS Management Console interface for EC2 instances. The left-hand navigation pane shows the 'Instances' tab selected. The main content area shows a table of instances. The instance 'KelsyCotto' (ID: i-00e1dc9cd6843478d) is highlighted in blue and is in a 'Stopped' state. A context menu is open over this instance, with the 'Start instance' option selected. The table below shows the details of the selected instance.

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv4 ...	Elast
Kcotto	i-0302a728fbd4095f	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-234-50-3.comp...	18.234.50.3	-
Instructor_test	i-088f1b7da8b8d0656	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-212-77-248.co...	18.212.77.248	-
KelsyCotto	i-00e1dc9cd6843478d	Stopped	m5.2xlarge	-	-	us-east-1a	-	-	-

Go to AWS EC2 Dashboard, select “Instances” tab, then find your instance. Right-click and chose ‘Start Instance’

When you restart your instance you will need to find your new Public DNS or IP address. Select your instance and “Connect” or look in Description tab. Then go back to instructions for “Logging into your instance”

The screenshot displays the AWS Management Console interface. On the left is a navigation sidebar with categories like EC2 Dashboard, Events, Tags, Limits, Instances, Images, Elastic Block Store, and Network & Security. The main area shows a list of EC2 instances. The instance 'KelsyCotto' is selected. Below the list, the 'Instance: i-00e1dc9cd6843478d (KelsyCotto)' details are shown, with tabs for Details, Security, Networking, Storage, Status checks, Monitoring, and Tags. The 'Details' tab is active, showing instance summary information. Two red arrows point to the 'Connect' button in the instance list and the 'Public IPv4 address' field in the instance details.

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv4 ...	Elast
Kcotto	i-0302a728fbd4095f	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-234-50-3.comp...	18.234.50.3	-
Instructor_test	i-088f1b7da8b8d0656	Running	m5.2xlarge	2/2 checks passed	No alarms	us-east-1a	ec2-18-212-77-248.co...	18.212.77.248	-
KelsyCotto	i-00e1dc9cd6843478d	Running	m5.2xlarge	Initializing	No alarms	us-east-1a	ec2-100-26-22-19.com...	100.26.22.19	-

Instance: i-00e1dc9cd6843478d (KelsyCotto)		
Details	Security	Networking
Instance summary		
Instance ID i-00e1dc9cd6843478d (KelsyCotto)	Public IPv4 address 100.26.22.19 open address	Private IPv4 addresses 172.31.28.82
IPv6 address -	Instance state Running	Public IPv4 DNS ec2-100-26-22-19.compute-1.amazonaws.com open address
Hostname type IP name: ip-172-31-28-82.ec2.internal	Private IP DNS name (IPv4 only) ip-172-31-28-82.ec2.internal	Elastic IP addresses -
Answer private resource DNS name IPv4 (A)	Instance type m5.2xlarge	

So, at this point:

- Your laptop/pc is ready for the workshop
- If it is not, you know where to get the information you need
- You know how to login to AWS
- The next step is to login to your linux machine on AWS and learn the basics of a linux command line

Break

Key AWS concepts and terminology

- **HDD** - Hard disk drive. A particular type of storage hardware that is generally cheaper and larger but slower than SSD. HDD drives are traditional hard drives that access data on a spinning magnetic disk.
- **Ephemeral storage** - Also known as Instance Store storage. Data storage associated with an EC2 instance that is local to the host computer. This storage does not persist when the instance is stopped or terminated. In other words, anything you store in this way will be lost if the system is stopped or terminated. Instance store volumes may be backed by SSD or HDD devices.

What is a Region?

- An AWS Region is set of compute resources that Amazon maintains (like the Data Center image shown before)
- Each Region corresponds to a physical warehouse of compute hardware (computers, storage, networking, etc.).
- At the time of writing there are 27 regions with more planned to come online soon: (US East (N.Virginia), US East (Ohio), US West (Oregon), US West (N. California), GovCloud (US-West), GovCloud (US-East), Canada (Central), EU (Ireland), EU (Frankfurt), EU (London), EU (Paris), EU (Milan), EU (Stockholm), Middle East (Bahrain), Middle East (UAE), Africa (Cape Town), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Seoul), Asia Pacific (Tokyo), Asia Pacific (Mumbai), Asia Pacific (Hong Kong), Asia Pacific (Beijing), Asia Pacific (Osaka), Asia Pacific (Jakarta), Asia Pacific (Ningxia), and South America (Sao Paulo).
- When you are logged into the AWS EC2 console, you are always operating in one of these regions.

What is a Region?

- Current region shown in the upper right corner of console
- It is important to pay attention to what region you are using for several reasons.
 - When you create an EC2 instance (EBS volume, etc) in one region you won't see it in another region.
 - The cost to use many AWS resources varies by region.
 - The region may influence network performance when you are accessing the instance, especially if you need to transfer large amounts of data in or out.
 - Billing is tracked separately for each region
 - Generally you should choose a region that is close to you or your users. But cost is also a consideration.